

Theoretical and methodological issues regarding the use of Language Technologies for patients with limited English proficiency

Harold Somers

School of Languages, Linguistics and Cultures
University of Manchester
Oxford Road, Manchester M13 9PL
Harold.Somers@manchester.ac.uk

Abstract

This paper concerns the use of spoken language translation as well as other technologies to support communication between clinicians and patients where the latter have limited proficiency in the majority language. The paper explores some theoretical and methodological issues, in particular the question of whether it is the patient or clinician who should be seen as the primary user of such software, and whether for certain scenarios more simple technology is preferable, especially given the huge overheads involved in developing SLT systems for under-resourced languages. A range of solutions are discussed.

1 Introduction

As its title suggests, this paper seeks to explore issues around the problem of using language technologies to support patients and healthcare providers where there is a significant language barrier. For convenience, in the title and elsewhere we use the phrase “patients with limited English proficiency (LEP)”, though it should be understood of course that much of the discussion would apply equally to other countries where the host or majority language is another language.

This paper is essentially theoretical and methodological, and although it does incorporate reflections on some recently completed pieces of research, it should be understood chiefly as a statement of the author’s views, and if it is in some respects confrontational or controversial, then this is in a sense deliberate.

In any western country there are recent or long-term immigrants, refugees, and asylum seekers and other people whose command of English, while often adequate for day-to-day activities such as shopping and other domestic chores, is not sufficient for more formal situations such as interactions with health services, especially visits to their doctor. There is no shortage of literature reporting disparities in healthcare provision in these communities and communication difficulties are identified as a major factor (e.g. Jones & Gill 1998, Fassil 2000, Jacobs et al. 2001, Bischoff et al. 2003, Flores et al. 2005, Westberg & Sorensen 2005), and an equally rich literature, which we will not review here, discusses traditional ways of addressing this problem, through use of interpreters and other services. Our interest is in the extent that language technology, including but not limited to machine translation (MT), may be able to provide some support as a contribution to a solution to this problem (Somers & Lovel 2003).

Two aspects of this issue need to be underlined immediately. First, it should be realised that this is a problem not just for the LEP patients, but for the healthcare providers with whom they need to interact: it is a matter not only of making oneself understood, but of understanding too. This seems to be an obvious point, but is often overlooked, for example in papers with titles referring to “problems of refugees” and so on, when more properly the focus should be on “problems of communication”. By the same token, note the use of the term “healthcare providers”: this is not just a problem for doctors, but for a wide range of professionals with whom patients must interact on the pathway to healthcare.

This brings us to the second point: while it is natural to focus on the doctor–patient consultation as the central element of the “pathway to healthcare”, in fact, this is only one of many diverse

interactions that a patient has with a variety of healthcare providers, including receptionists at clinics and hospitals, paramedics, nurses, therapists, pharmacists as well, of course, as the “doctor” who may be a GP, a consultant, a specialist, and so on. Each of these interactions involves a range of communicative activities requiring different language skills and implying different language technologies, often but not inevitably involving *translation* in some form.

In this paper we will first explore this issue of different users and different scenarios, always focusing on how particular aspects of this impact on the choice and design of language technology.

We will then look in particular at the doctor–patient interview and compare the relatively sophisticated approach of using Spoken Language Translation (SLT) as compared to use of much simpler technology, as tested in some recent research by the present author.

2 Different users, different scenarios

As stated above, although it is natural to think of “going to the doctor” as involving chiefly an interview with a GP, and while everything in medical practice arguably derives from this consultation, the pathway to healthcare in normal circumstances involves several other processes, all of which involve language-based encounters that present a barrier to LEP patients. Let us consider the range of processes, interlocutors, and possible technologies that might be suitable, reiterating some points made previously by this author (Somers 2006).

2.1 The pathway to healthcare

The pathway might begin with a person suspecting that there may be something wrong with them. Many people nowadays would in this situation first try to find out something about their condition on their own, typically on the Word-Wide Web. If you need this information in your own language, and you have limited literacy skills, as is the case with many asylum seekers and refugees, technologies implied are multilingual information extraction, MT perhaps coupled with text simplification, with synthesized speech output. For specific conditions which may be treated at specialist clinics (our own experience is based on Somalis with respiratory difficulties) it may be possible to identify a series

of frequently asked questions and set up a pre-consultation computer-mediated help-desk and interview (cf. Osman et al. 1994).

Having decided that a visit to the doctor is indicated, the next step is to make an appointment. Appointment scheduling is the classical application of SLT, as seen in most of the early work in the field, and is a typical case of a task-oriented cooperative dialogue. Note that the dialogue partner – the receptionist in the clinic – does not necessarily have any medical expertise, nor possibly the high level of education and openness to new technology that is often assumed in the literature on SLT.

If this is the patient’s first encounter with this particular healthcare institution, they may wish to get their “history”, a task nowadays often done separately from the main doctor–patient consultation, to save the doctor’s time. This might be a suitable application for computer-based interviewing (cf. Bachman 2003).

The next step might be the doctor–patient consultation itself, which has been the focus of much attention (e.g. papers at the recent *Workshop on Medical Speech Translation* at HLT/NAACL in New York in 2006). While some developers (e.g. Bouillon et al. 2005) originally assumed that the patient’s role in this can be reduced to simple responses involving yes/no responses, gestures and perhaps a limited vocabulary of simple answers, current clinical theory in contrast focuses on *patient-centred* medicine (cf. Stewart et al. 2003), an approach now adopted by Bouillon et al. (2007). The session will see the doctor eliciting information in order to make a diagnosis as foreseen, but also explaining the condition and the treatment, exploring the patient’s feelings about the situation, and inviting the patient to ask questions. So the dialogue is very much a two-way interaction. Of course this presents massive difficulties for SLT system design.

After the initial consultation, the next step may involve a trip to the pharmacist to get some drugs or equipment. Apart from the human interaction, the drugs (or whatever) will include written instructions and information: frequency and amount of use, contraindications, warnings and so on. This is an obvious application for controlled language MT: drug dose instructions are of the same order of complexity as weather bulletins, though there remains the practical problem of

transferring the text from the packet to the translation system. For non-literate patients, “talking pill boxes” are already available (marketed by MedivoxRx, see Orlovsky 2005), so it would be nice if they could “talk” in a variety of languages.

Another outcome might involve another practitioner – a nurse or a therapist – and a series of meetings where the condition may be treated or managed. Apart from more scheduling, this will almost certainly involve explanations and demonstrations by the practitioner, and typically also elicitation of further information from the patient. Hospital treatment would involve interaction with a wide range of staff, again not all medical experts.

All this introduces the question of who is the principle user of a communication device, which will have a bearing on many design issues. In contrast for example with several medical SLT designs, where it is assumed that the doctor is the one who controls the dialogue and accordingly controls the SLT system interface (Narayanan et al. 2004:101, Bouillon et al. 2005, Ettelaie et al. 2005:89), we might propose that it is the *patient* who is going to be the regular user, and who should therefore “own” the device.

At the very least, it should be recognised that a communication device (whether SLT or some other technology, see below) will typically have two users at any time, who may have very different skills and expectations, and these need to be taken into consideration in the design. Indeed, just like the healthcare providers, as already mentioned, not all patients are alike, and they may represent a wide range of levels of language ability (both native and target), literacy, computer literacy, and a variety of expectations and experiences regarding healthcare itself. It is therefore obvious that interfaces to any communication systems should be flexible, and possibly different depending on the profile of the user.

Realistically, we are not going to address all these problems, but let us consider some of the basic technology issues that the different usage scenarios introduce.

2.2 Language technology implications


Our discussion so far has mentioned or implied a number of basic technologies including SLT, text MT, multilingual information extraction, text

simplification, and computer-based interviewing, automatic speech recognition (ASR) and speech synthesis. Let us focus on applications involving translation.

One obvious problem for these technologies is that often the language we are interested is one of the so-called “under-resourced” languages: this severely limits what can be done, and precludes for example using off-the-shelf components, since they simply do not exist. The effort required to develop SLT for an under-resourced language should not be underestimated (cf. Black et al. 2002, Schultz et al. 2004, Zhou et al. 2004, Narayanan et al. 2004, 2006, Kathol et al. 2005, Besacier et al. 2006, Schultz & Black 2006). We have explored the possibility of “faking” speech synthesis as an interim solution to this (Evans et al. 2002, Somers et al. 2006) with a fairly promising evaluation based on the doctor–patient dialogue scenario using a German synthesizer to produce fake Somali output. Currently we are attempting the more audacious task of “fake” speech recognition by tricking an English ASR system into recognizing a limited vocabulary of Urdu words, with astonishingly good results when the system has to choose from a set of possible responses (Rizvi, in prep.).

Even with languages that are better resourced, developing applications suitable for this scenario can be challenging. For example, Wang (2007) reports a Chinese–English SLT system built by pipelining commercially available Chinese ASR, Chinese–English MT and English speech synthesis, tested once again in the healthcare scenario. Replicating the evaluation methodology of Somers & Sugita (2003) in which subjects are asked to identify the intended meaning of a translated answer to a specific question, he found that Chinese ASR is the weakest link in the chain with around 70% correct interpretation of ASR+MT, dropping to 62% when output is synthesized. MT on its own was 97% understandable. This differs from the finding reported in Somers & Sugita (2003), where Japanese ASR was quite reliable, and MT was the weak link. Chinese ASR is evidently considerably more difficult.

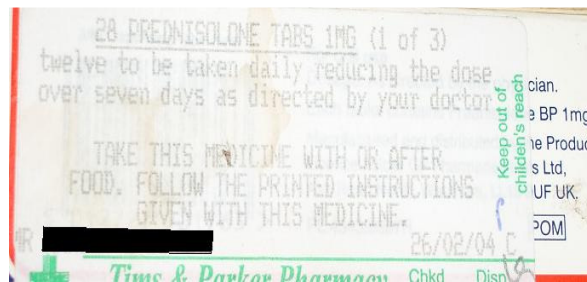
Taking ASR out of the equation still requires text to be input. Exploring the scenario of LEP patients wishing to read prescription labels, Ghobadi (2007) first experimented with a handheld



```

~R Cai.Dms 0.1s
put ~dr~ into the nostrils
tWiCE jalnis directedr'
THROB QNAY OGYS N~IS-
OPENLY NOT aII@.

```



```

is. PREDN~SOL@E Teaks i~'r Al of 3)
t~eliio, -be taken Baily, reducing ile Diane
;er seven days as dire~liied bj yc'ur ~ucIjI-S
li THIS ~E ~ AFTER
e. Fa_lo* Tie )~.~IEU INITRU~TIO~i
~"~ 'ijIj'~i%::"

```

Figure 1. Images of typical prescription labels and results of scanning with a handheld scanner.

scanner. If one considers that a typical prescription label is printed with a low-quality printer onto a label that is then often wrapped round a container, it is no surprise that scan results leave a lot of work still to be done (cf. Figure 1).

If instructions cannot be scanned in, we must devise some other text input method suitable for a user who does not know English, may not be familiar with the Roman script, and may be illiterate, even in their own language. The obvious solution is to have the labels translated at source, i.e. by the pharmacist, though this involves huge problems related to the pharmacist's legal obligation to verify the instructions on the label, which obviously they cannot do if they are written in a foreign language. Despite some political opposition, LEP is recognized in the US as a potential source of discrimination, and a 1998 Office of Civil Rights memorandum (OCR 1998) puts in place requirements for translations to be made available as part of healthcare provision. There is some evidence of use of MT (e.g. Sharif et al. 2006, Barclay 2007) where available, which of course always needs to be checked for translation accuracy, but this is not a viable solution for many of the languages needed. And even where the foreign language in question (Spanish) is well resourced, there is a reluctance to do so (Barclay 2007).

3 Spoken Language Translation vs. low-level technology

The problem of LEP patients has had some attention from the Language Technology

community: so far, the focus has been on medical SLT systems, as mentioned above. We have elsewhere (Somers 2006) made some critical remarks about the direction some of this research has taken, and these are worth briefly repeating here in connection with our proposal that SLT – especially as currently implemented – is not always the most appropriate technology for all LEP patients' (and their clinicians') needs.

We have already mentioned the fact that current SLT systems inevitably see the doctor as being in control of the system and hence of the dialogue itself. Several assumptions underlying this set-up are false: the doctor's familiarity with computers in general and the SLT device in particular is assumed to be superior to the patient's (e.g. Narayanan et al. 2004:101, Precoda et al. 2004:9, Bouillon et al. 2007:42), but this may not be true, especially if the patient becomes a regular user. In our own research, admittedly with a much simpler device (Johnson 2007, Somers & Lovel 2007, Somers et al. in prep.), we found many patients more than willing to share or even take over control of the device, as shown in Figure 2a, in contrast to the scenario presented in the on-line video demo of an SLT system (Figure 2b), where the doctor (the one in the white coat) has total control to the extent that the patient is not even allowed to see the screen.

Sharing the device will also facilitate its use in promoting communication via a combination of technologies. Text and (where literacy is a problem) pictures can support the spoken (translated) word and even to a certain extent supplant speech: certain parts of doctor-patient



(a) Clinician and patient sharing the laptop device (from Somers & Lovel 2007)



(b) Snapshot from Transonics' demo movie (source: <http://sail.usc.edu/transonics/demo/transedit02lr.mov>, accessed 14 May 2007)

Figure 2. Contrasting perspectives in use of computer-based communication device by clinician and patient

dialogues (and indeed other exchanges on the pathway to healthcare) follow a fairly predictable pattern that can be exploited by using predetermined questions and (sets of) possible answers which, as we have discovered (Johnson 2007, Somers & Lovel 2007, Somers et al. in prep.) can lead to very high satisfaction rates, even though some drawbacks are recognized

In our research, in which as a test case we focused on Somalis with asthma-related conditions, we developed software on an ordinary laptop using a mousepad, and on a touch-screen tablet using a stylus, which permitted clinicians to choose freely from a range of 69 questions grouped under various topics. The questions were presented in both English and Somali, with pre-recorded (human) speech for both the questions and the possible answers on a screen as illustrated in Figure 3. The patient could review all the possible answers by clicking on the symbols before indicating to the clinician the desired answer.

We tested the software in simulated consultations with six GPs and asthma nurses and 26 Somali patients. All 26 simulations were completed adequately: none were abandoned due to difficulties using the system, with communication, or due to frustration on the part of Somalis or clinicians. In 20 of the 26 simulations, all questions were answered by the patients. Post-session feedback questionnaires indicated extremely high satisfaction ratings by both clinicians and patients with almost every aspect of the system (see Table 1): the only serious drawback noted was the rather obvious problem

that the system did not allow the users to go off-script, as reflected in low clinician satisfaction scores for eliciting the patient's worries (42%) and building a relationship (69%), both key contributors to the overall goal of achieving a clinical outcome (65%).

Of course the system described does not involve MT in any sense. The reason for mentioning it here is to make the point that for some aspects of doctor-patient communication, where the content of the dialogue is sufficiently predictable, it might be safer to use a simpler technology such as that described here. We will surely need SLT for some communication tasks, but it makes sense, especially when the effort



Figure 3. Screen showing possible answers to the question “What kind of animal did you own in Somalia?”. The question itself, and each of the possible answers is associated with a digitised recording in Somali.

	VS	S	other	%
Size of symbols (P)	25	1	0	100
Size of symbols (C) N=9	5	4	0	100
Size of text (P)	23	1	2	92
Size of text (C) N=9	4	5	0	100
Range of questions (P)	25	1	0	100
Range of questions (C) N=9	1	7	1	89
Range of responses (P)	21	3	2	92
Range of responses (C) N=9	3	5	1	89
Using laptop (P) N=14	3	3	8*	43
Using tablet (P) N=12	7	4	1	91
Using mousepad (P) N=14	3	3	8*	43
Using stylus (P) N=12	9	3	0	100
Navigation (P)	11	9	6*	77
Navigation (C)	14	12	0	100
P's ability to use device (P)	8	12	6*	77
P's ability to use device (C)	12	9	5*	81
C's ability to use device (P)	26	0	0	100
C's ability to use device (C)	9	15	0	100
P understand C's questions	23	3	0	100
C understand P's responses	10	13	3	88
P answer C's questions	22	4	0	100
C elicit information	12	11	3	88
Make self understood (P)	22	4	0	100
Make self understood (C)	8	15	3	88
P explain worries to C	11	4	1	96
C elicit P's worries	7	4	13	42
Build a relationship (P)	22	1	3	88
Build a relationship (C)	7	11	8	69
Better than no interp. (P)	22	3	1	96
Better than no interp. (C)	14	8	4	85
P satisfied with review	25	1	0	100
C achieved desired outcome	11	6	9	65

Table 1. Satisfaction ratings for a variety of questions. Key: "P" patient (N=26), "C" clinician (N=9). "VS" very satisfied, "S" satisfied, "Other" includes dissatisfied, very dissatisfied, don't know and (especially where marked *) not applicable. Except where indicated, N=26, corresponding to the number of sessions.

required to build SLT systems for certain languages is so great, to seek alternative solutions.

4 Conclusions

Spoken language translation and MT for under-resourced languages are two of greatest new challenges for the MT community. Putting them together gives a task that is almost impossible to contemplate at the present time. In this paper we have looked at one particular domain where the need for such technology is particularly important, and, in the spirit of the title of the TMI conference

series, have put forward some theoretical and methodological issues related to that task. The main theoretical point made has been the need to focus on user-centered rather than technology-centered design in SLT. And regarding methodology, the point has been made that some lesser technologies, as well as some "cheats", may be the way forward, at least in the short term.

References

- Bachman JW (2003) The patient-computer interview: a neglected tool that can aid the clinician. *Mayo Clinic Proceedings* 78, 67–78.
- Barclay L (2007) Pharmacies may not always translate prescription labels for non-English speaking patients. [Report on presentation by L Weiss at Society for General Internal Medicine 2007 Annual Meeting, Toronto, Ont. Abstract 172022]. *Medscape Medical News*, April 27, 2007. Available at: <http://www.medscape.com/viewarticle/555840>. Accessed 14 May 2007.
- Besacier L, Le V-B, Boitet C, Berment V (2006) ASR and translation for under-resourced languages. *IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, Toulouse, France, pp. V-1221–4.
- Bischoff A, Bovier PA, Isah R, Françoise G, Ariel E, Louis L (2003) Language barriers between nurses and asylum seekers: their impact on symptom reporting and referral. *Social Science in Medicine* 57, 503–12.
- Black AW, Brown RD, Frederking RF, Singh R, Moody J, Steinbrecher E (2002) TONGUES: Rapid development of a speech-to-speech translation system. *Proceedings of HLT 2002: Second International Conference on Human Language Technology Research*, San Diego, California, pp. 183–6.
- Bouillon P, Flores G, Starlander M, Chatzichrisafis N, Santaholma M, Tsourakis N, Rayner M, Hockey BA (2007) A bidirectional grammar-based medical speech translator. *Proceedings of the Workshop on Grammar-based approaches to spoken language processing*, Prague, Czech Republic, pp. 41–48.
- Bouillon P, Rayner M, Chatzichrisafis N, Hockey BA, Santaholma M, Starlander M, Nakao Y, Kanzaki K, Isahara H (2005) A generic multi-lingual open source platform for limited-domain medical speech translation. *Proceedings of the Tenth Conference on European Association of Machine Translation*,

- "Practical applications of machine translation", Budapest, Hungary, pp. 50–8.
- Ehsani F, Kimzey J, Master D, Park H, Sudre K (2006) Rapid development of a speech translation system for Korean. *IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, Toulouse, France, pp. V-1225–8.
- Ettelaie E, Gandhe S, Georgiou P, Belvin R, Knight K, Marcu D, Narayanan S, Traum D (2005) Transonics: A practical speech-to-speech translator for English-Farsi medical dialogues. *43rd Annual Meeting of the Association for Computational Linguistics: ACL-05 Interactive Poster and Demonstration Sessions*, Ann Arbor, MI, pp. 89–92.
- Evans G, Polyzoaki K, Blenkhorn P (2002) An approach to producing new languages for talking applications for use by blind people. In K Miesenberger, J Klaus, W Zagler (eds) *8th ICCHP, Computers Helping People with Special Needs*, (LNCS 2398), Berlin: Springer Verlag, pp. 575–82.
- Fassil Y (2002) Looking after the health of refugees. *British Medical Journal* 321, 59.
- Flores G, Abreu M, Tomany-Korman SC (2005) Limited English proficiency, primary language at home, and disparities in children's health care: How language barriers are measured matters. *Public Health Report* 120, 418–30.
- Ghobadi B (2007) Farsi prescription labels translator. Dissertation, School of Informatics, University of Manchester.
- Jacobs EA, Lauderdale DS, Meltzer D, Shorey JM, Levison W, Thisted R (2001) Impact of interpreter services on delivery of health care to limited-English-proficient patients. *Journal of General Internal Medicine* 16, 468–74.
- Johnson MJ (2007) An exploration into support for communication between health care practitioners and Somalis using assistive language technology in the context of asthma consultations. PhD thesis, School of Nursing, Midwifery and Social Work/School of Informatics, University of Manchester.
- Jones D, Gill P (1998) Refugees and primary care: tackling the inequalities. *British Medical Journal* 317, 1444–6.
- Kathol A, Precoda K, Vergyri D, Wang W, Riehemann S (2005) Speech translation for low-resource languages: The case of Pashto. *Interspeech'2005 – Eurospeech*, Lisbon, pp. 2273–6.
- Narayanan S., Ananthakrishnan S, Belvin R, Ettelaie E, Gandhe S, Ganjavi S, Georgiou PG, Hein CM, Kadambe S, Knight K, Marcu D, Neely HE, Srinivasamurthy N, Traum, D, Wang D (2004) The Transonics spoken dialogue translator: An aid for English-Persian doctor-patient interviews. In T Bickmore (ed.) *Dialogue Systems for Health Communication: Papers from the 2004 Fall Symposium*, Menlo Park, California: American Association for Artificial Intelligence, pp. 97–103.
- Narayanan SS, Georgiou PG, Sethy A, Wang D, Bulut M, Sundaram S, Ettelaie E, Ananthakrishnan S, Franco H, Precoda K, Vergyri D, Zheng J, Wang W, Gadde RR, Graciarena M, Abrash V, Frandsen M, Richey C (2006) Speech recognition engineering issues in speech to speech translation system design for low resource languages and domains. *IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, Toulouse, France, pp. V-1209–12.
- OCR (Office of Civil Rights) (1998) Guidance memorandum: Title VI prohibition against national origin discrimination—persons with limited-English proficiency. January 29, 1998. Available at: <http://www.hhs.gov/ocr/lepfinal.htm>. Accessed 14 May 2007.
- Orlovsky C (2005) Talking pill bottles let medications speak for themselves'. *NurseZone.com* (online magazine). Available at: www.nursezone.com/Job/DevicesandTechnology.asp?articleID=14396. Accessed 14 May 2007.
- Osman L, Abdalla M, Beattie J, Ross S, Russell I, Friend J, Legge J, Douglas J (1994) Reducing hospital admissions through computer supported education for asthma patients. *British Medical Journal* 308, 568–71.
- Precoda K, Franco H, Dost A, Frandsen M, Fry J, Kathol A, Richey C, Riehemann S, Vergyri D, Zheng J, Culy C (2004) Limited-domain speech-to-speech translation between English and Pashto. *HLT-NAACL 2004, Human Language Technology Conference of the North American Chapter of the Association for Computational Linguistics: Demonstrations*, Boston, MA, pp. 9–12.
- Rizvi SM (in prep.) Faking Urdu speech recognition for a doctor-patient dialogue system. MSc dissertation, School of Informatics, University of Manchester.
- Schultz T, Alexander D, Black AW, Peterson K, Suebisai S, Waibel A (2004) A Thai speech translation system for medical dialogs. *HLT-NAACL 2004, Human Language Technology Conference of the North American Chapter of the Association for Computational Linguistics: Demonstrations*, Boston, MA, pp. 34–5.

- Schultz T, Black AW (2006) Challenges with rapid adaptation of speech translation systems to new language pairs. *IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, Toulouse, France, pp. V-1213–16.
- Sharif I, Lo S, Ozuah PO (2006) Availability of Spanish prescription labels. *Journal of Health Care for the Poor and Underserved* 17, 65–9.
- Somers H (2006) Language Engineering and the pathway to healthcare: A user-oriented view. *HLT/NAACL-06 Medical Speech Translation, Proceedings of the Workshop*, New York, pp. 32–9.
- Somers HL, Caress A-L, Evans DG, Johnson MJ, Lovel HJ, Mohamed Z (in prep.) A computer-based aid for communication between patients with limited English and their clinicians, using symbols and digitised speech. Submitted to *International Journal of Medical Informatics*.
- Somers H, Evans G, Mohamed Z (2006) Developing speech synthesis for under-resourced languages by “faking it”: An experiment with Somali. *5th International Conference on Language Resources and Evaluation*, Genoa, Italy, pp. 2578–81.
- Somers H, Lovel H (2003) Computer-based support for patients with limited English. *Association for Computational Linguistics EACL 2003, 10th Conference of the European Chapter, Proceedings of the 7th International EAMT Workshop on MT and other language technology tools*, Budapest, pp. 41–9.
- Somers HL, Lovel HJ (2007) Can AAC technology facilitate communication for patients with limited English? ESRC Project Final Report, School of Informatics, University of Manchester. Available at <http://www.informatics.manchester.ac.uk/~harold/ESRCfinal.pdf>. Accessed 8 August 2007.
- Somers H, Sugita Y (2003) Evaluating commercial spoken language translation software. *MT Summit IX: Proceedings of the Ninth Machine Translation Summit*, New Orleans, pp. 370–7.
- Stewart M, Brown JB, Weston WW, McWhinney IR, McWilliam CL, Freeman TR (2003) *Patient-centered medicine: Transforming the clinical method* (2nd ed.). Abingdon, Oxon: Radcliffe.
- Wang B (2007) Chinese to English speech translation system built from standard components. Dissertation, School of Informatics, University of Manchester.
- Westberg SM, Sorensen TD (2005) Pharmacy-related health disparities experienced by non-English-speaking patients: Impact of pharmaceutical care. *Journal of the American Pharmaceutical Association* 45, 48–54.
- Zhou B, Déchelotte D, Gao Y (2004) Two-way speech-to-speech translation on handheld devices. *INTERSPEECH 2004 – ICSLP, 8th International Conference on Spoken Language Processing*, Jeju Island, Korea, pp. 1637–40.